Local anisotropy analysis for non-smooth images

Sandra Bergonnier, François Hild, Stéphane Roux

► To cite this version:

Sandra Bergonnier, François Hild, Stéphane Roux. Local anisotropy analysis for non-smooth images. Pattern Recognition, 2007, 40 (2), pp.544-556. hal-00014052

HAL Id: hal-00014052

<https://hal.science/hal-00014052>

Submitted on 17 Dec 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Local anisotropy analysis for non-smooth images

Sandra BERGONNIER ^{a,b} François HILD ^a and Stéphane ROUX ^{b,1}

^a*LMT-Cachan*

ENS de Cachan / CNRS-UMR 8535 / Université Paris 6

61 avenue du Président Wilson, F-94235 Cachan Cedex, France.

Courriel: Sandra.Bergonnier@lmt.ens-cachan.fr, Francois.Hild@lmt.ens-cachan.fr

^b*Laboratoire “Surface du Verre et Interfaces”,*

UMR CNRS/Saint-Gobain 125

39 quai Lucien Lefranc, 93303 Aubervilliers Cedex

Courriel: stephane.roux@saint-gobain.com

Abstract

Identification of local anisotropy and determination of principal axes is addressed through different methods that are designed to be tolerant to the non-smooth character of images on the pixel scale. These different tools are validated on various examples, and their performances are compared. The most powerful, robust and accurate method consists in computing the curvature tensor of the auto-correlation function of regularized images using fast Fourier transforms.

Key words: Texture analysis, padding, orientation, fibrous material.

PACS:

1 Introduction

Estimating the anisotropy of images or zones of images, and the corresponding orientation field has a wide range of applications in texture analysis. From

¹ Correspondant. Tel: +33 1 48 39 57 52, Fax: +33 1 48 34 74 16.

fingerprints to transmission electron microscopy, many cases require the identification of anisotropy as a major characteristic of texture or pattern recognition. As a result, numerous techniques for estimating local orientation fields have been proposed in the past.

One of the earliest approaches, introduced by Rao [1], to determine the principal orientation field from an image uses the direction of local gray-level gradients. The local principal anisotropy direction is perpendicular to the gradient. Hanbury et al. [2] used the algorithm suggested by Rao to describe the vector field of oriented textures. From this first analysis, mathematical morphology was applied to characterize the texture, namely circular centered gradient operator for texture segmentation and circular centered top-hat operator for defect detection. The method was validated [2] through the analysis of a wood plank picture, *i.e.*, segmentation into regions of uniform orientation.

Since then, a number of variants have been introduced, most of them also based on gradients. One such popular method is based on the Principal Component Analysis (PCA) of local gradient [3,4]. By applying the PCA to the autocovariance matrix of the gradient vectors provides the 2-dimensional probability density function of these vectors. The main direction of gradient is then obtained from the autocovariance matrix. Determining the local orientation can be used for segmentation or defect detection. Bazen et al. [5] used PCA applied to the autocovariance matrix of the gradient vectors for fingerprint segmentation. First the image is low-pass filtered for noise suppression, then the directional field is estimated and finally a coherence-based segmentation provides the identification of the basic shape of the fingerprint used for classification, encoding and recognition.

Some improvements on methods first based on the calculation of gradients have been developed. Gu et al. [6] proposed a new approach for the orientation field of fingerprints. It is based on the combination of the two models by a weight function. A polynomial model is used to approximate the orientation field globally and a point-charge model at each singular point is used to improve the approximation locally. Feng et al. [7,8] presented a technique based on the Principal Component Analysis and the multiscale decomposition to improve the robustness to noise. Stuke et al. [9] proposed a method to analyze superimposed oriented patterns. The definition of a generalized structure tensor (GST) is used to detect the number of orientation based on confidence measures calculated by the invariants of the GST.

All the above methods rely on the strong hypothesis that gradients are well-defined, and hence that they can be computed and used in a safe way. Another very successful method has been developed by Bigün et al. [10–12] for images (movies) that are invariant in one (or more) directions. The latter provides an anisotropy orientation computed in the least squares sense in the Fourier

domain. Some more technical details will be presented in Section 2.

Often associated with the determination of the orientation field, the estimation of the amplitude of the local anisotropy is an important issue. Germain et al. [13] proposed a scattering indicator that associates anisotropy measurements with the scale at which the texture is observed. Scharcanski et al. [14] first used an extension of Rao’s algorithm to determine the orientation and then proposed a measure for the local anisotropy considering that the stronger the local anisotropy, the higher the local alignment of gradient vectors.

Most of the established techniques are hence based on the analysis of the local gradient field, which make them very sensitive to noise. Schematically summarizing most of the above mentioned approaches, different filtering procedures are introduced to restore a meaningful gradient field, from which the basic algorithm can be used. In a number of well-documented cases, these approaches provide a satisfactory tool, which may be tolerant to small amplitude uncorrelated noise. However, there remain cases, such as the glass wool images studied in a further section of the present paper, where these approaches give poor results. The key question is to know whether it is legitimate or not to filter the image without losing information, or in other words, if noise and information-carrying parts of the image can be easily separated through an adapted filtering.

The purpose of this paper is to present and discuss a different philosophy of the approach from which different algorithms are proposed. A class of methods that *do not* require the underlying image to be smooth (at least in the sense of differentiability so that a meaningful gradient can be computed) is considered. The paper is organized as follows. Section 2 presents a brief description of some classical tools alluded to above. In Section 3, the key requirement of robustness and sensitivity is introduced, together with general considerations on the class of non-smooth textures that may be addressed. In Section 4, different methods are presented and applied to some test cases in Section 5. Their relative performances are compared. An objectivity measurement is proposed to validate the robustness of the method as a function of the quality of the analyzed image.

2 Principles of anisotropy analysis

Most of the existing methods are based on the calculation of a gradient field. Let $f(x, y)$ represent a gray-scale image, and $\nabla f(x, y)$ its gradient. Based on the discrete (pixel-based) field, a “suited” approximant $\mathbf{G}$ — this term is kept as vague as possible since numerous choices of the latter have been designed to extend the robustness and accuracy of the measured anisotropy — of the

gradient vector ∇f is computed and denoted as

$$\mathbf{G} = \begin{pmatrix} G_x(x, y) \\ G_y(x, y) \end{pmatrix} \quad (1)$$

Rao [1] proposes to compute over windows W the tensor

$$\mathbf{\Gamma} = \sum_W \mathbf{G} \otimes \mathbf{G} \quad (2)$$

where $\otimes$ denotes the tensorial product. The direction of the anisotropy axis is then perpendicular to the eigenvector associated to the largest eigenvalue.

Bazen et al. [4] defined the autocovariance matrix as the diadic tensor product of gradient vectors averaged over a zone of interest, *i.e.*, the tensor $\mathbf{\Gamma}$ defined in Eq. (2). The main direction of the gradient is then given by the eigenvector of the autocovariance matrix $\mathbf{\Gamma}$ that corresponds to the largest eigenvalue. A local *strength* of the directional field that corresponds to an anisotropy degree is defined as

$$Str = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} \quad (3)$$

where λ_1 and λ_2 correspond to the largest and smallest eigenvalues of the autocovariance matrix $\mathbf{\Gamma}$ (Eq. (2)), respectively.

Bigün et al. [10–12] were interested in strongly anisotropic images or movies, (such as optical flow conserving images) displaying a strict invariance along a space(-time) direction. This invariance implies that the power spectrum is concentrated over a plane (hyperplane) normal to the invariance direction. Thus the orientation axis is determined thanks to a minimization of a moment of inertia with respect to an axis going to the origin, in Fourier space. The power density $|\hat{f}(k)|^2$ is interpreted as a mass density, and thus the inertia tensor $\mathbf{T}$ reads

$$\mathbf{T}_{ij} = \iint_{E_2} k_i k_j |\hat{f}(k)|^2 dk_1 dk_2 \quad (4)$$

The eigenvector of $\mathbf{T}$ associated to the smallest eigenvalue gives the anisotropy direction.

3 Preliminary remarks concerning anisotropy analysis

Detecting the anisotropy of a domain, hereafter called Region Of Interest or ROI, contained in the image involves an evaluation of the change of gray levels in different directions, the direction where the variation is the least being the principal axis of anisotropy. Since *changes* in gray levels are to be considered, most methods are based on gradients. However, although it appears to be an

unavoidable statement, the very definition of gradients on discrete maps raises the key aspect of pixel scale sensitivity of tools based on further exploitation of discrete gradient fields. Different aspects may even be listed at this level of generality:

- *Scale sensitivity.* Anisotropy amplitude and direction are not intrinsic concepts but dependent on the scale of analysis, as discussed in details in various articles addressing either extraction of anisotropy direction or estimate of anisotropy amplitude [13].
- *Robustness of the analysis.* Random noise is ubiquitous. The robustness of the method will thus significantly depend on the explicit or implicit choice of the privileged scale of analysis. For instance in order to reduce the pixel scale sensitivity, some methods use a Sobel gradient operator that allows one to smooth out part of it.
- *Smoothness of the image.* Not unrelated to the previous property, it is very often considered that in contrast to noise that is essentially assumed to be spatially uncorrelated, the reference image is (a discretization of) a smooth function, with very low power content in high frequency modes. In other words, the reference image is assumed to be “differentiable”, in the sense that gradient estimates using different scale estimators will give a quasi-identical value for a minimum range of scales from one to, say, ten pixels. Let us note that this can be an intrinsic property of the original image to be analyzed, or it can also result from a low-pass filtering of the image that will restore smoothness for a regularity-lacking image. However in the latter case, it is important to ensure that this filtering does not erase the spatial frequency component responsible for the dominant anisotropy.

As a result of the previous considerations, most methods implicitly assume a smoothness of the image that allows for the use of gradient estimates. The further extraction of the anisotropy axis then consists in seeking for a direction that is orthogonal (on average) to the gradient field. Methods that are based on a different strategy are proposed, namely, one would like to avoid the *a priori* introduction of any smoothness assumption. This may be an important improvement for images where the anisotropy is provided by very high frequency modes (*i.e.*, fine texture). In the following section three different methods are presented.

4 Proposed methods

4.1 Power-based method

To illustrate the discussion of Section 3, Fig. 1 (top) shows an example of an anisotropic texture, which contains high frequency contributions. In fact, for this simple example many wavelengths are responsible for the same anisotropy. However, as shown in Fig. 1 (bottom), through cuts along the x or y directions going through the center of the image, there are very drastic changes of gray levels at one pixel distance. Therefore in this case, gradients will depend strongly on the discrete operator used to estimate them. This map has been obtained from a self-affine function [15] that lacks differentiability in the continuum limit.

Even if the notion of a gradient as an intrinsic quantity independent of the operator fails, from the cuts along the x and y directions, meaningful differences appear markedly. The lower curve in Fig. 1 (bottom), $f(x, y_c)$ at fixed $y = y_c$, displays very rapid variations at the pixel scale, whereas the upper one, $f(x_c, y)$ at fixed $x = x_c$ is much smoother at short wavelengths, while still displaying longer wavelength components. This suggests the use of an L_2 -norm of the finite difference along the $\mathbf{e}$ direction

$$\|D_{\mathbf{e}}f\| = \left[\sum_{(x,y)} (f(x + e_x, y + e_y) - f(x, y))^2 \right]^{1/2} \quad (5)$$

where e_x, e_y are the components of the vector $\mathbf{e}$. Taking the above example, such a norm applied only to the cuts shown in Fig. 1 (bottom), gives respectively 29 and 11 gray-levels for $f(x, y_c)$ and $f(x_c, y)$. It is noteworthy that the use of this norm does not rely on any smoothness assumption for f .

This example allows one to clarify the key property one aims to use in order to extract the anisotropy axis. In classical approaches, once noise is properly extracted, the signal is assumed to be smooth, and thus on a local scale, one may observe patterns as shown in Fig. (2) (top), where a mean gradient is easily detected, hence providing an estimate of the anisotropy axis as perpendicular to the gradient direction. In contrast, non-smooth images where typical local pattern may rather appear as schematically depicted in Fig. (2) (bottom) is addressed. There, instead of estimating the gradient, which may not be well defined, one may focus on the total power, or L_2 -norm of the subimage in different directions.

4.2 Polar mapping

One possibility is to extract an annulus (radius r , $R_0 < r < R_1$) centered at any point of interest (Fig. 3(a)) and to map it onto a rectangle $[R_0, R_1] \times [0, 2\pi]$ using polar coordinates (Fig. 3(b)). This mapping requires a subpixel interpolation of the gray levels that can be conveniently performed using Q1 elementary functions $(1, x, y, xy)$ over the natural square pixel array. Unfolding the annulus onto a rectangle, one can compute the above norm over radial differences for each angle. A sequence $g(\theta) = \|D_{e_r} f(r, \theta)\|$ is obtained as a 2π -periodic function. Based on the fact that this norm will vary as a function of the relative orientation of the radial vector, and the anisotropy axis, it is expected that $g(\theta)$ display a $\cos(2(\theta - \theta_0))$ component proportional to the anisotropy (Fig. 3(c)).

To evaluate the anisotropy axis, this function is projected onto the $e^{2i\theta}$ function, namely computing the real and imaginary parts

$$\begin{aligned} S_g &= (1/N) \sum_{j=1}^N \sin(2\theta_j) g(\theta_j) \\ C_g &= (1/N) \sum_{j=1}^N \cos(2\theta_j) g(\theta_j) \end{aligned} \tag{6}$$

where the index j runs over the N discrete orientations of the 2π interval, and $\theta_j = 2j\pi/N$. The amplitude of the anisotropy A is then given by the norm of this Fourier component, $A = \sqrt{C_g^2 + S_g^2}$, whereas the anisotropy direction is given by the phase ϕ such that $C_g + iS_g = -Ae^{2i\phi}$.

The dual to this approach can also be considered through the computation of the L_2 -norm along the orthoradial (θ) direction, still averaged over radial directions, $h(\theta) = \|D_{e_\theta} f(r, \theta)\|$. This function is expected now to be maximum at the orientation of the anisotropy axis. Again, h is decomposed along the sine and cosine of 2θ

$$\begin{aligned} S_h &= (1/N) \sum_{j=1}^N \sin(2\theta_j) h(\theta_j) \\ C_h &= (1/N) \sum_{j=1}^N \cos(2\theta_j) h(\theta_j) \end{aligned} \tag{7}$$

The amplitude of this vector also provides an evaluation of the anisotropy amplitude, $A = \sqrt{C_h^2 + S_h^2}$, whereas the anisotropy direction is given by the phase ϕ such that $C_h + iS_h = Ae^{2i\phi}$.

The time-consuming step of this algorithm is the mapping of the annulus onto a rectangle domain. An alternative still based on the L_2 -Norm is proposed in the following subsection.

4.3 Anisotropy tensor

Let us note that the L_2 -norm can in fact easily be computed for any direction in Fourier space. Indeed, the total power is preserved in the Fourier transform, and thus

$$\|D_{\mathbf{e}}f\|^2 = \sum_{\mathbf{k}} (\mathbf{k} \cdot \mathbf{e})^2 |\hat{f}(\mathbf{k})|^2 \quad (8)$$

Therefore, introducing the anisotropy tensor $\mathbf{T}$ defined in Eq. (4) one observes that

$$\|D_{\mathbf{e}}f\|^2 = \mathbf{T} : (\mathbf{e} \otimes \mathbf{e}) \quad (9)$$

where ‘:’ is the tensorial contraction with respect to two indices. Therefore, the anisotropy axis appears as the minor eigendirection of the anisotropy tensor $\mathbf{T}$. The anisotropy amplitude can be estimated conveniently through $\text{dev}(\mathbf{T})/\text{tr}(\mathbf{T})$, where tr designates the trace, and dev the norm of the deviator $\mathbf{T} - (1/2)\text{tr}(\mathbf{T})\mathbf{I}$. Let us note that up to a factor 2, this definition would be similar to the above defined “strength”, Str , in Eq. (3) computed over a different tensor.

The latter formulation avoids the unnecessary and time-consuming polar mapping, and gives quite naturally access to the anisotropy after a numerically efficient Fast Fourier Transform (FFT) step. This scheme is here identical to the method introduced by Bigün et al. [10–12]. Let us note however that the routes followed to reach this result were very different.

4.4 Auto-correlation method

Yet another approach is to compute the auto-correlation function, $C(\mathbf{x})$ of the considered region within the image

$$C(\mathbf{x}) = \iint f(\mathbf{y})f(\mathbf{y} + \mathbf{x})d\mathbf{y} - \left(\iint f(\mathbf{y})d\mathbf{y}\right)^2 \quad (10)$$

The way this function decays with respect to the distance to the origin $\mathbf{r} = 0$ represents the progressive loss of similarity of the pattern as one moves along a given direction. Therefore, it constitutes a good candidate to estimate the anisotropy. Moreover, the auto-correlation function is easily computed through FFT algorithms with a very moderate computational cost. To quantify the anisotropy and find the preferential orientation of the texture, the autocorrelation function is truncated at a fixed level, chosen in this study to be 80% of the maximum value $C(0)$. This allows us to define the domain $\mathcal{D}$ where $C(r)/C(0) > 0.8$. This domain is then characterized by its geometrical inertia tensor

$$\mathbf{I} = \iint_{\mathcal{D}} \mathbf{y} \otimes \mathbf{y} d\mathbf{y} \quad (11)$$

This tensor is finally diagonalized, and its principal axis (direction of the largest eigenvalue) is the anisotropy axis. The amplitude of anisotropy can again be estimated using the expression of the strength, Eq. (3), or $\text{dev}(\mathbf{I})/\text{tr}(\mathbf{I})$.

4.5 Connection between the two last proposals

The two last methods, as presented, appear as quite different. However, they are in fact connected. As the threshold $C(r)/C(0)$ is chosen close to unity, the autocorrelation function for a smooth image can be approximated by a paraboloid. Therefore the domain $\mathcal{D}$ will tend toward an ellipse. The inertia tensor of this ellipse $\mathbf{I}$ can thus simply be related to the curvature of the autocorrelation function at the origin. Now, the latter curvature is easily computed in Fourier space, and is identical to the $\mathbf{T}$ tensor introduced above. Therefore $\mathbf{T}^{-1}$ is proportional to $\mathbf{I}$ when the threshold $C(r)/C(0)$ tends to unity. One can easily check that the expression of the anisotropy amplitude computed over $\mathbf{T}$ or $\mathbf{I}$ is identical.

4.6 Non smooth images

Up to now, the non-smooth character of the image has not been considered. However, based on the last comment, it is straightforward to see how to extend the present analysis to such non smooth images. The non-smoothness can be characterized in a quantitative (statistical) fashion by a Holder index [15], ζ , such that

$$\langle [f(\mathbf{x} + \mathbf{y}) - f(\mathbf{x})]^2 \rangle \sim |\mathbf{y}|^{2\zeta} \quad (12)$$

The ζ index characterizes the maximum order of differentiability of f . Because it is a discrete field, it is always possible to form finite difference “gradients”, however as mentioned above, the latter will suffer from a strong sensitivity to the local weight used for such an estimate. In terms of the auto-correlation function, a singularity appears at the origin $C(\mathbf{r})/C(0) = 1 - A|\mathbf{r}|^{2\zeta}$, and thus the above referenced “parabolic” approximation gives a poor picture of the correlation for $\zeta < 1$. The curvature tensor, $\mathbf{T}$, is also controlled by the pixel size. Let us note that the case $\zeta = 1$ corresponds to a differentiable or smooth function.

The main drawback of encountering a cases where $\zeta < 1$, is the fact that the small scale controls the estimators, and hence the sensitivity to noise becomes very detrimental to their quality. In order to compensate for this, one may resort to a suited filtering that makes gradients well defined. The natural way of carrying out this operation is to perform an integration of order $1 - \zeta$. In real space this implies a convolution of the original image with

a power-law kernel. In Fourier space, it simply consists in multiplying the Fourier transform $\hat{f}(k)$ by $|k|^{\zeta-1}$. This operation transforms any image into a smooth one, where gradients can now be defined safely without the pixel scale sensitivity mentioned before.

The use of such a filter for the present method is almost transparent. In terms of the anisotropy tensor, one simply has to compute

$$\mathbf{T}_\zeta = \iint_{E_2} \frac{\mathbf{k} \otimes \mathbf{k}}{|\mathbf{k}|^{2-2\zeta}} |\hat{f}(k)|^2 d\mathbf{k} \quad (13)$$

From the latter tensor, the subsequent treatment remains similar to the previous analysis.

For the autocorrelation method, since $C(r)$ is computed in Fourier space, it suffices to multiply by the corresponding power-law of k before reverting to real space, and thresholding the corresponding autocorrelation function. By construction, in the vicinity of the origin, the parabolic approximation is then suited.

Finally, for the present procedure to be operational, the index ζ has to be known. The latter is a signature of the statistical texture of the image. To evaluate ζ , the simplest way is to study the average over the different regions of interest, ROI, of the power spectra, and to fit it with a power-law of $|k|$. The exponent of the power-law $-\beta$ is related to the index ζ through $\beta = 2(1 + \zeta)$. This defines operationally the procedure where ζ is no longer a parameter to be manually adjusted, but rather set through a well-defined characterization of the considered image. Thus no free parameter remains.

4.7 Neutral padding

One important issue concerning the use of the present tool to analyze a texture is the size of the region of interest (ROI). The latter is bounded by two different properties. On the one hand, an upper limit is dictated by the texture itself. The anisotropy axis should not vary too significantly within one single ROI, and thus its size should be smaller than the correlation length of the orientation field. Note that this upper bound is a physical scale, and not related to the pixel size. On the other hand, the size of the ROI cannot be shrunk to too few pixel sizes. One reason is that the information content has to be rich enough. But the main reason is more subtle, and results from the use of Fourier transforms. The latter technique assumes implicitly that the analyzed field is periodic, and hence the right and left hand sides are adjacent to each other, as well as the bottom and top rows. In general, this artificial assumption leads to strong discontinuities along the border that can easily be misinterpreted as fixing

the texture dominant orientation. Let us note that this is an edge effect, and thus its weight as compared to the bulk vanishes as the ROI size increases. In practice, for the studied images, ROI sizes less than 32×32 pixels gave unreliable results. (Note that FFT calls for an integer power of 2 as ROI sizes, so that less than 32×32 pixels means less than or equal to 16×16 pixels).

To reduce the lower bound, we have developed an original procedure, which allows us to deal with 16×16 pixel or even 8×8 pixel ROIs quite safely. It consists in padding the image in such a way that a periodicity is enforced without altering the original image in the core of the ROI. Let us assume that one wishes to work with an $N \times N$ ROI size. For convenience in the present discussion, we assume the origin to be placed at the image center, $-N/2 < x < N/2$ and $-N/2 < y < N/2$. We first embed the image in a larger $M \times M$ frame with the same center. The difference between the two is a domain which will be used to implement the periodicity, through a suited padding, called “neutral padding”. In order not to introduce an additional information that may perturb the image, the high frequency components in Fourier space are set to zero. The number of modes is chosen to be equal to $M^2 - N^2$, so that the problem is uniquely defined. As such, the solution to this problem can easily be formulated as a simple linear system to solve. However, being partly formulated in real space, and partly in Fourier space, makes its solution difficult to write explicitly in a simple manner. This would thus call for a long computation of the solution. However, an *exact* solution to the problem is not needed. Thus a simple iterative scheme (fixed point formulation) is introduced and stops after a few iterations, so that an approximate answer is obtained.

More precisely, let us call $f(x, y)$ the original image, and $h^{(m)}(x, y)$ the m th iterate of the padded image. One starts from a first guess $h^{(0)}(x, y) = f(x, y)$ for $-N/2 < x < N/2$ and $-N/2 < y < N/2$ (see Figure 4(a)). The outskirts of the images from the core to the border is filled by any available data; either the original image, if it is defined, or an average gray value. The image is then filtered in Fourier space $\hat{h}^{(1)}(k_x, k_y) = \hat{h}^{(0)}(k_x, k_y)$ for $-N/2 < k_x < N/2$ and $-N/2 < k_y < N/2$. Outside this core region, $\hat{h}^{(1)}(k_x, k_y) = 0$. Reverting to real space, one obtains a function $h^{(1)}(x, y)$ that now differs from f in its core region. These values $h^{(2)}(x, y) = f(x, y)$ are forced back if $-N/2 < x < N/2$ and $-N/2 < y < N/2$, otherwise $h^{(2)}(x, y) = h^{(1)}(k_x, k_y)$. The procedure from $h^{(0)}$ to $h^{(2)}$ defines one step to be iterated. Upon iterations, h converges toward the fixed point that maps exactly the prescribed image in the core, and zero Fourier amplitude in the high frequency domain. In practice 10 iterations provide a good approximation to the exact solution, because the convergence is exponential (see Figure 4(b)). By construction, M has to be an integer power of 2. However N can assume any value strictly less than M . To validate the impact of this procedure, $N = M/2$ is chosen.

For the application of the methods, the neutral padding can be used equivalently either on the original image or on the image first treated with the integration filter.

5 Application to test images

5.1 *Fingerprint*

The above presented methods have been applied to different images. The first image to be tested is fingerprints where ridges form the oriented texture to analyze. The difficulty comes from the interruption of the ridges, and the rotation of the anisotropy close to the central vertex. The image size is 240×240 pixels, partitioned in ROIs of size 32×32 pixels. These regions are centered on a regular square grid with half overlap between the zones.

The results shown in Fig. 5 are the two orientations of the polar L_2 -method, the auto-correlation technique and the anisotropy tensor approach. The ratio between the two radii of the annulus for the polar method is 2, and the angle increment is 2° . For the polar method, results from the radial and orthoradial orientation are quite similar, and also comparable to those obtained by the two other methods apart from the sites located close to the edge where the signal is poor. However, at those sites, the auto-correlation and the anisotropy tensor methods appear as more satisfactory. One significant difference yet is the computation time that appears to be much smaller (about 100 times) for the auto-correlation and anisotropy tensor techniques as compared to the polar L_2 method. As previously mentioned this is due to the rectangular to polar mapping that is a consuming step of the algorithm.

5.2 *A non-smooth image: wood*

If the non-smooth character of the image is considered, the original anisotropy tensor method can be compared with the same method coupled with the integration filter. The picture of a wood plank is a good candidate for a non-smooth image. The sequence of tree rings and the knot can be a complex application for the original method. Figure 6(top) shows the results of the algorithm using the anisotropy tensor method, Fig. 6(middle) the results when the integration filter is added to the method and Fig. 6(bottom) the results when the neutral padding is used after the integration filter. There are slight differences between the results although all the methods provide reliable estimates even if no known reference exists. For Fig. 6(top) the orientations

globally correspond to the visually estimated angles except for the right part of the plank core. In Fig. 6(middle) the plank core is better solved and in Fig. 6(bottom) both the orientations of the plank core and the knot have a better concordance with the visually estimated angles.

5.3 Concentric circles

To test the performances of the different methods an example where the true orientation is known at each position in the image is needed. This is achieved by studying a picture made of concentric circles that represent an absolute reference. The image size is 512×512 pixels, and the size of the ROIs 32×32 pixels. These regions are centered on a regular square grid and the distance between the centers of two consecutive ROIs is 16 pixels.

In Fig. 7, where only one quarter of the images is shown, the two orientations of the polar L_2 -method, the auto-correlation technique and the anisotropy tensor approach are presented. A first visual inspection gives a rough evaluation of the adequation between real orientations and estimates. For the four methods, the calculation points are the same. The size of the region of interest is shown on the lower right corner of all images. For the two polar methods, the results are quite similar and in good agreement with the true orientation. The directional field determined using the anisotropy tensor also corresponds accurately to the expected result. Concerning the auto-correlation method, the results are quite poor. The orientation does not fit with the circles. This can be explained by the form of the auto-correlation function for this particular image. When the auto-correlation function is truncated at a fixed level, the resulting domain has to be simply connected. In Fig. 8 the auto-correlation function and its projection after truncation are represented. In that case, the projection consists in several domains. That method is not suited to this particular image. It would be straightforward to mend the method so that only the connected component to the origin is extracted. This however requires a computation time that disqualifies the method.

Figure 9 shows the differences between the determined and true orientations as a function of the angle for the three valid methods. The 0° angle corresponds to a horizontal orientation. When the anisotropy tensor method with the integration filter and the neutral padding is considered, the differences with the true orientation are quite small with an RMS error equal to 0.35° . It is to be noted that the integration filter and the neutral padding allows one to gain one order of magnitude for the RMS error compared to the raw anisotropy tensor method. The results given by the polar radial and orthoradial methods are the best compared to the true orientation. The RMS error is equal to 0.053° for the L_2 radial method and to 0.046° for the L_2 orthoradial

method with a computation time equal to 178 s and 184 s respectively. This example reveals that there is no method that outperforms the other ones in terms of the estimated orientation as for the three methods the RMS error is smaller than 1° , but the computation time appears to be smaller (about 3 to 4 times) for the anisotropy tensor method with the integration filter and the neutral padding.

5.4 Crimped glasswool

The method can also be tested on a more complex application, namely determination of the texture orientation of a crimped glass wool sample. For such images, it is to be noticed that there is a very fine texture, with a very difficult partition between noise and information carrying signals. Moreover, if anisotropy is clearly marked, the coherence length is rather short, and the pattern appears as circumvoluted. On such images, classical gradient based methods appear as inappropriate. This is thus a very severe test of the two other methods.

Figure 11(top) shows the results of the polar L_2 orthoradial method and Fig. 11(middle) the results of the anisotropy tensor method with the integration filter to treat the non-smooth character of the image and the neutral padding, particularly suited to the detection of very fine texture. If an absolute measure of the quality of the method cannot be assessed on real images, one can propose a quantitative comparison between the last two methods. However, such a comparison should accommodate for the π -periodicity of the angular determination. For a given ROI centered on $\mathbf{x}$, the difference between the angles determined by both methods, the polar L_2 orthoradial method and the complete anisotropy tensor method, is calculated. Figure 11(bottom) shows the squared anisotropy amplitude of the complete anisotropy tensor method as a function of the angle difference between the two methods. The widespread values is a mere reflection on the difficulty to estimate these orientations. One can however note that large angle differences occur more frequently when the anisotropy amplitude is low (*i.e.*, for a low anisotropy).

5.5 Objectivity measurement

For a severe test case of the method, such as the crimped glass wool image, the objectivity of the method is tested, *i.e.*, the property that the orientation is independent of the magnification/resolution of the image. From a glass wool image, whose original size is 512×512 pixels, the resolution of the image has been artificially degraded, *i.e.*, this operation leads to the construction of images at different scales. Scale no. 0 corresponds to the considered image. From

scale no. 1 on, each transition is characterized by the definition of super-pixels. The latter are defined recursively from one scale to the next by averaging the gray levels of 2×2 neighboring pixels. This procedure is carried out until the minimum size of the sub-image is equal to 128 pixels. The size of the ROIs is divided by 4 from one scale to the next. The size of the real image is 512×512 pixels and the ROI covers 32×32 pixels. The size of the scale no. 1 picture will be 256×256 and the size of the ROI 16×16 pixels. For the last scale, corresponding to a 128×128 pixels image, the size of the ROI will be 8×8 pixels.

The technique used for this objectivity measurement is the anisotropy tensor method with integration filter and neutral padding. Figure 12 shows the results for the three analyzed scales. The three maps correspond to the square cosine of the angle values of each ROI. The results on scale no. 0 are given in Fig. 12(top), on scale no. 1 in Fig. 12(middle) and on scale no. 2 in Fig. 12(bottom). An excellent agreement is observed on the first two maps, in spite of the fact that 16 pixels is a very small ROI size. On scale no. 2 the coarsening leads to significant discrepancies. Let us underline that the corresponding ROI size, of 8 pixels is extremely small and induces a severe sensitivity to the preferential axes.

6 Conclusions

Three methods to estimate the local orientation of a texture have been described and tested on different images whose textures are complex. Although inspired by a significantly different route as traditional approaches, the resulting tools are in some respect closely related. The results given by the methods have then been compared in terms of accuracy and computation time:

- In terms of accuracy in anisotropy orientation, on examples where the texture is known, the auto-correlation method fails to provide a reasonable estimate. The other methods give a good result, with a slightly more accurate result for the polar L_2 methods. On examples where the texture is more complex, the polar and the complete anisotropy tensor methods give rather different results for the same size of analyzed regions, but large angle differences occur when the anisotropy is not clearly marked. The complete anisotropy tensor method can detect very fine textures.
- Moreover, it has been shown that the obtained result is objective, and that it can handle poor resolution even with very small regions of interest (down to 16 pixel size).
- In terms of computation time, performances of the autocorrelation and the anisotropy tensor methods are much better than the polar and angular L_2 -methods.

The conclusion is that the complete anisotropy tensor method outperforms all others for real images. Let us note that in the case of smooth images, this method reduces to the one introduced by Bigün et al. [10,11]. The extension to non-smooth images allows one to deal safely with commonly encountered cases, without introducing any free parameter. Moreover the extension to neutral padding allows one to deal with very fine textures using small size of analyzing window with a reliable estimate of the orientation independently of the resolution of the tested image.

7 Acknowledgments

This work was performed within a collaboration with Isover Saint-Gobain. The authors also wish to thank M. Baudequin, C. Machelart et J.-B. Rieunier for useful discussions.

About the author – SANDRA BERGONNIER received an engineering degree at Institut Français de Mécanique Avancée in 2002. Currently, she is a PhD student working on the mechanical characterization of entangled materials.

About the author – FRANÇOIS HILD received his MS degree in Mechanical Engineering from the University of Paris 6 in 1989. He received his PhD degrees in Mechanical Engineering from the University of Paris 6 in 1992 and from the University of California in 1995. Currently, he is research professor (CNRS) and works at LMT-Cachan. His research interests include optical techniques and texture characterization in Mechanics of solids and structures.

About the author – STEPHANE ROUX received an engineering degree at Ecole Polytechnique and Ecole Nationale des Ponts et Chaussées, before a doctorate degree at Ecole Supérieure de Physique et Chimie Industrielles de Paris. Currently, he is a research director at the CNRS, and director of a joint research laboratory with the CNRS (French public research organisation) and Saint-Gobain (an international material manufacturer company). His research interests cover statistical physics and solid mechanics, focusing in particular on the behaviour of heterogeneous materials.

References

- [1] A. Rao, *A Taxonomy for Texture Description and Identification*, Springer-Verlag, 1990.
- [2] A. Hanbury, J. Serra, Analysis of oriented textures using mathematical morphology, in: *Vision with Non-Traditionnal Sensors*, 2002.
- [3] A. Bazen, S. Gerez, Directional field computation for fingerprints based on the principal component analysis of local gradients, in: *ProRISC 2000 Workshop on Circuits, Systems and Signal Processing*, 2000.
- [4] A. Bazen, G. Verwaaijen, S. Gerez, L. Veelenturf, B. van der Zwaag, A correlation-based fingerprint verification system, in: *ProRISC 2000 Workshop on Circuits, Systems and Signal Processing*, 2000.
- [5] A. Bazen, S. Gerez, Systematic methods for the computation of the directional field and singular points of fingerprints, *EEE Transactions on Pattern Analysis and Machine Intelligence* Volume 24.
- [6] J. Gu, J. Zhou, D. Zhang, A combination model for orientation field of fingerprints, *Pattern Recognition* Volume 36 (2003) 775–990.
- [7] X. Feng, P. Milanfar, Multiscale principal components analysis for image localorientation estimation, in: *Proceeding of the 36th Asilomar Conference on Signals, Systems and Computers*, 2002.
- [8] X. Feng, Analysis and approaches to image local orientation estimation, Master’s thesis, UC Santa Cruz (2003).
- [9] I. Stuke, T. Aach, E. Barth, C. Mota, Analysing superimposed oriented patterns, in: *6th IEEE Southwest Symposium on Image Analysis and Interpretation*, 2004, pp. 133–137.
- [10] J. Bigun, G. Granlund, Optimal orientation detection of linear symmetry, in: *Proceeding of the IEEE first international conference on computer vision*, 1987.
- [11] J. Bigun, G. Granlund, J. Wiklund, Multidimensional orientation estimation with applications to texture analysis and optical flow, *IEEE-PAMI* Volume 13 (1991) 775–990.
- [12] K. Nilsson, J. Bigun, Localization of corresponding points in fingerprints by complex filtering, *Pattern Recognition Letters* Volume 24 (2003) 2135–2144.
- [13] C. Germain, J. D. Costa, O. Lavialle, P. Baylou, Multiscale estimation of vector field anisotropy application to texture characterization, *Signal Processing* Volume 83 (2003) 1487–1503.
- [14] J. Scharcanski, C. Dodson, Stochastic texture image estimators for local spatial anisotropy and its variability, *IEEE Transactions on Instrumentation and Measurement* Volume 49.
- [15] J. Feder, *Fractals*, Plenum Press, New-York, 1988.

List of Figures

- 1 (top) An example of an anisotropic picture, where smoothness is insufficient to allow for a proper gradient definition.
(bottom) Gray levels along horizontal and vertical cuts through the previous image. Note that although the rapid variation on the pixel scale does not allow for a proper gradient definition, the power spectra of the two cuts are very different. 20

- 2 Schematic plot of the gray level pattern extracted in order to use (top) standard gradient based approach, (bottom) non-smooth pattern as in the present approach. 21

- 3 (a) considered annulus (radius r , $8 < r < 16$ pixels), (b) its polar mapping with an angular increment $= 2^\circ$, (c) the L_2 -norm along the radial direction as a function of θ (solid line) and its $\cos(2(\theta - \theta_0))$ component (dashed line). 22

- 4 (a) original ROI: $h^{(0)}(x, y) = f(x, y)$ prior to neutral padding ($N = 32$ pixels), (b) neutral padding on the ROI: result after 10 iterations ($M = 64$ pixels). 23

- 5 (top left) Polar L_2 radial method (external radius $= 32$ pixels, radius ratio $= 2$, angle increment $= 2^\circ$) computation time: 36.9 s. (top right) Polar L_2 orthoradial method (external radius $= 32$ pixels, radius ratio $= 2$, angle increment $= 2^\circ$) computation time: 33.4 s. (bottom left) Auto-correlation method: (threshold $C(r)/C(0)=0.8$) computation time: 1.10 s. (bottom right) Anisotropy tensor method: computation time: 0.33 s. 24

- 6 Comparison of the results on a wood plank (ROI size $= 32 \times 32$ pixels, distance between two ROI centers $= 16$ pixels) (top) original angle map obtained with the anisotropy tensor method, (middle) resulting angle map with the integration filter, (bottom) neutral padding added to the integration filter. 25

7	(top left) Polar L_2 radial method (external radius = 32 pixels, radius ratio = 2, angle increment = 2°) computation time: 178 s, (top right) Polar L_2 orthoradial method (external radius = 32 pixels, radius ratio = 2, angle increment = 2°) computation time: 184 s, (bottom left) Auto-correlation method (threshold $C(r)/C(0) = 0.8$) computation time: 4.85 s, (bottom right) Anisotropy tensor method with the integration filter and the neutral padding, computation time: 53.4 s.	26
8	(top) Auto-correlation function of one region of interest. (bottom) Projection of the truncated auto-correlation function.	27
9	Differences between calculated and true orientations for the polar L_2 radial method (top), the polar L_2 orthoradial method (middle) and the anisotropy tensor method with the integration filter and the neutral padding (bottom).	28
10	Robustness of the method to noise: (left) Differences between a noise-safe image (left half) and a noisy image with a relative standard-deviation of 35.5% (right half). (right) Average error angle in degree as a function of the relative noise standard-deviation.	29
11	(top) Results of the L_2 orthoradial method on a glass wool image (external radius = 32 pixels, ratio = 2, angle increment = 2°), distance between two ROI centers = 16 pixels), (middle) Results of the anisotropy tensor method on a glass wool image (ROI size = 32×32 pixels, distance between two ROI centers = 16 pixels), (bottom) Squared anisotropy amplitude of the anisotropy tensor method as a function of the angle differences between the auto-correlation and the anisotropy tensor methods.	30
12	Square cosine of the orientation map for (top) a 512×512 pixel glass wool image, and a 32×32 pixel ROI, (middle) a 256×256 pixel reduced image, and a 16×16 pixel ROI, (bottom) a 128×128 pixel reduced image, and a 8×8 pixel ROI.	31

Fig. 1. (top) An example of an anisotropic picture, where smoothness is insufficient to allow for a proper gradient definition. (bottom) Gray levels along horizontal and vertical cuts through the previous image. Note that although the rapid variation on the pixel scale does not allow for a proper gradient definition, the power spectra of the two cuts are very different.

Bergonnier et al.

Fig. 2. Schematic plot of the gray level pattern extracted in order to use (top) standard gradient based approach, (bottom) non-smooth pattern as in the present approach.

Bergonnier et al.

Fig. 3. (a) considered annulus (radius r , $8 < r < 16$ pixels), (b) its polar mapping with an angular increment $= 2^\circ$, (c) the L_2 -norm along the radial direction as a function of θ (solid line) and its $\cos(2(\theta - \theta_0))$ component (dashed line).

Bergonnier et al.

-a-

-b-

Fig. 4. (a) original ROI: $h^{(0)}(x, y) = f(x, y)$ prior to neutral padding ($N = 32$ pixels), (b) neutral padding on the ROI: result after 10 iterations ($M = 64$ pixels).

Bergonnier et al.

Fig. 5. (top left) Polar L_2 radial method (external radius = 32 pixels, radius ratio = 2, angle increment = 2°) computation time: 36.9 s. (top right) Polar L_2 orthoradial method (external radius = 32 pixels, radius ratio = 2, angle increment = 2°) computation time: 33.4 s. (bottom left) Auto-correlation method: (threshold $C(r)/C(0)=0.8$) computation time: 1.10 s. (bottom right) Anisotropy tensor method: computation time: 0.33 s.

Bergonnier et al.

Fig. 6. Comparison of the results on a wood plank (ROI size = 32×32 pixels, distance between two ROI centers = 16 pixels) (top) original angle map obtained with the anisotropy tensor method, (middle) resulting angle map with the integration filter, (bottom) neutral padding added to the integration filter.

Bergonnier et al.

Fig. 7. (top left) Polar L_2 radial method (external radius = 32 pixels, radius ratio = 2, angle increment = 2°) computation time: 178 s, (top right) Polar L_2 orthoradial method (external radius = 32 pixels, radius ratio = 2, angle increment = 2°) computation time: 184 s, (bottom left) Auto-correlation method (threshold $C(r)/C(0) = 0.8$) computation time: 4.85 s, (bottom right) Anisotropy tensor method with the integration filter and the neutral padding, computation time: 53.4 s.

Bergonnier et al.

Fig. 8. (top) Auto-correlation function of one region of interest. (bottom) Projection of the truncated auto-correlation function.

Bergonnier et al.

Fig. 9. Differences between calculated and true orientations for the polar L_2 radial method (top), the polar L_2 orthoradial method (middle) and the anisotropy tensor method with the integration filter and the neutral padding (bottom).

Bergonnier et al.

Fig. 10. Robustness of the method to noise: (left) Differences between a noise-safe image (left half) and a noisy image with a relative standard-deviation of 35.5% (right half). (right) Average error angle in degree as a function of the relative noise standard-deviation.

Bergonnier et al.

Fig. 11. (top) Results of the L_2 orthoradial method on a glass wool image (external radius = 32 pixels, ratio = 2, angle increment = 2°), distance between two ROI centers = 16 pixels), (middle) Results of the anisotropy tensor method on a glass wool image (ROI size = 32×32 pixels, distance between two ROI centers = 16 pixels), (bottom) Squared anisotropy amplitude of the anisotropy tensor method as a function of the angle differences between the auto-correlation and the anisotropy tensor methods.

Bergonnier et al.

Fig. 12. Square cosine of the orientation map for (top) a 512×512 pixel glass wool image, and a 32×32 pixel ROI, (middle) a 256×256 pixel reduced image, and a 16×16 pixel ROI, (bottom) a 128×128 pixel reduced image, and a 8×8 pixel ROI.

Bergonnier et al.